

**UNITED STATES PATENT APPLICATION**

**of**

**Ross KONINGSTEIN**

**Valentin SPITKOVSKY**

**Georges R. HARIK**

**and**

**Noam SHAZEER**

**for**

**SUGGESTING AND/OR PROVIDING TARGETING  
CRITERIA FOR ADVERTISEMENTS**

**Expr ss Mail No. EL997188685US**

## **SUGGESTING AND/OR PROVIDING TARGETING CRITERIA FOR ADVERTISEMENTS**

### **§ 1. BACKGROUND OF THE INVENTION**

5

#### **§ 1.1 FIELD OF THE INVENTION**

The present invention concerns advertising. In particular, the present invention concerns improving keyword targeting used to serve ads.

#### **10 § 1.2 BACKGROUND INFORMATION**

Advertising using traditional media, such as television, radio, newspapers and magazines, is well known. Unfortunately, even when armed with demographic studies and entirely reasonable assumptions about the typical audience of various media outlets, advertisers recognize that much of their ad budget is simply wasted. Moreover, it is very difficult to identify and eliminate such waste.

Recently, advertising over more interactive media has become popular. For example, as the number of people using the Internet has exploded, advertisers have come to appreciate media and services offered over the Internet as a potentially powerful way to advertise.

Advertisers have developed several strategies in an attempt to maximize the value of such advertising. In one strategy, advertisers use popular presences or means for providing interactive media or services (referred to as "Websites" in the specification without loss of generality) as conduits to reach a large audience. Using this first approach, an advertiser may place ads on the home page of the New York Times Website, or the USA Today Website, for example. In another strategy, an advertiser may attempt to target its ads to more narrow niche audiences, thereby increasing the likelihood of a positive response by the audience. For example, an agency promoting tourism in the Costa Rican rainforest might place ads on the ecotourism-travel subdirectory of the Yahoo Website.

Regardless of the strategy, Website-based ads (also referred to as “Web ads”) are typically presented to their advertising audience in the form “banner ads” (i.e., a rectangular box that may include graphic components). When a member of the advertising audience (referred to as a “viewer” or “user” in the

5 Specification without loss of generality) selects one of these banner ads by clicking on it, embedded hypertext links typically direct the viewer to the advertiser’s Website. This process, wherein the viewer selects an ad, is commonly referred to as a “click-through” (“Click-through” is intended to cover any user selection.).

10 Advertisers may judge the efficacy of an advertising campaign using a number of measurable or determinable user behaviors, such as click-throughs, click-through rates, conversions, conversion rates, etc. The ratio of the number of click-throughs to the number of impressions of the ad (i.e., the number of times an ad is displayed) is commonly referred to as the “click-through rate” of the ad.

15 A “conversion” is said to occur when a user consummates a transaction related to a previously served ad. What constitutes a conversion may vary from case to case and can be determined in a variety of ways. For example, it may be the case that a conversion occurs when a user clicks on an ad, is referred to the advertiser’s Web page, and consummates a purchase there before leaving that

20 Web page. Alternatively, a conversion may be defined as a user being shown an ad, and making a purchase on the advertiser’s Web page within a predetermined time (e.g., seven days). In yet another alternative, a conversion may be defined by an advertiser to be any measurable/observable user action such as, for example, downloading a white paper, navigating to at least a given depth of a

25 Website, viewing at least a certain number of Web pages, spending at least a predetermined amount of time on a Website or Web page, registering on a Website, etc. Often, if user actions don’t indicate a consummated purchase, they may indicate a sales lead, although user actions constituting a conversion are not limited to this. Indeed, many other definitions of what constitutes a conversion

30 are possible. The ratio of the number of conversions to the number of impressions of the ad (i.e., the number of times an ad is displayed or otherwise

rendered) is commonly referred to as the conversion rate. If a conversion is defined to be able to occur within a predetermined time since the serving of an ad, one possible definition of the conversion rate might only consider ads that have been served more than the predetermined time in the past.

5       Despite the initial promise of Website-based advertisement, there remain several problems with existing approaches. Although advertisers are able to reach a large audience, they are frequently dissatisfied with the return on their advertisement investment. Some have attempted to improve ad performance by tracking the online habits of users, but this approach has led to privacy concerns.

10

### **§ 1.2.1       TARGETTED AD SERVING**

A popular recent trend has been to target ads to users based on some type of user request, such as the submission of a search query to a search engine. For example, the Google search engine Website allows advertisers to 15 specify keywords for triggering the serving of an ad or a group of ads when those keywords, or some derivative thereof, are included in a search query. Unfortunately, it can often be difficult for advertisers to specify appropriate keywords, or some other targeting criteria, for a given ad. Such poor targeting may lead to inappropriate ad serves (“over-serving”), or no ad servers when 20 appropriate (“under-serving”).

#### **§ 1.2.1.1       OVER-SERVING AND CONSEQUENT PROBLEMS**

In ad serving systems in which advertisers pay per selection, rather than 25 per impression, most advertisers are not troubled if they get non-relevant impressions. After all, if users don't select the ads, the impressions are free. Some Website hosts have chosen to place advertising revenues over the interests of users. One such Website is “Overture.com”, which hosts a so-called 30 “search engine” service returning advertisements masquerading as “search results” in response to user queries. The Overture.com website permits advertisers to pay to position an ad for their Website (or a target Website) higher up on the list of purported search results. If schemes where the advertiser only

pays if a user selects the ad (i.e., cost-per-click) are implemented, the advertiser lacks incentive to target their ads effectively, since a poorly targeted ad will not be selected and therefore will not require payment. Consequently, high cost-per-click ads show up near or at the top, but do not necessarily translate into real

5 revenue for the ad publisher because viewers don't click on them. Furthermore, ads that users might otherwise select may be further down the list, or not on the list at all, and so relevancy of ads is compromised. Accordingly, such systems are certainly not good from the standpoint of users seeking relevant information, since their interests are subordinated to those of the advertisers. However, when  
10 such systems employ a cost per click payment scheme rather than a cost per impression scheme, they are not particularly efficient from the standpoint of generating revenue for the ad server.

Some ad serving systems, such as Adwords Select from Google for example, will shut off low click-through keywords, and account for the performance of ads in serving decisions. In such systems, advertisers may need to pick relevant keywords for their ads.

Further, sometimes targeting criteria will trigger the serving of an ad at an inappropriate or undesirable time. For example, although a keyword targeting criteria (e.g., “travel”) of an ad (e.g., an ad for business travel) may literally match a term of a user’s search query (e.g., “space travel”), serving the ad (e.g., an ad for business travel) might sometimes be inappropriate or undesirable. This could affect the performance of the ad. In some advertising systems, poor ad performance can result in increased costs and/or the ad being dropped.

## 25                   § 1.2.1.2    UNDER-SERVING AND CONSEQUENT                           PROBLEMS

In the case where an advertiser did not consider a particular targeting criteria for its ad, there could be a lost opportunity to serve an otherwise relevant ad for presentation to an end user because of the missing targeting criteria. For example, an advertiser selling "Star Trek" DVDs may not think of using "Captain Kirk" as a targeting criteria, although such a query would likely be indicative of a user with an interest in "Star Trek" movies. This example illustrates problems

associated with missed opportunities, where an ad isn't served when doing so would be appropriate or desirable.

From the perspective of the entity serving ads, missed opportunities to serve a relevant ad are lost opportunities to generate revenue. For example,

5 with some ad serving systems, such as Adwords Select from Google for example, many search result pages served may contain no (or few) ads. If such opportunities to show an ad are missed, not only are they lost forever as a source of revenue, but nothing is learned about what ads could be shown for these searches, consequently forsaking future revenue.

10 From the perspective of an advertiser, not only does the advertiser miss an opportunity to serve its ad, but if its ad were served pursuant to a match with a less popular keyword, the cost of such an ad serve may be less expensive than a more popular keyword with more competing advertisers.

15 **§ 1.2.1.3 SELECTING EFFECTIVE SERVING CRITERIA SUCH AS KEYWORDS**

As the sections on over-serving and under-serving above illustrate, poorly targeted ads are bad for users, advertisers, and ad serving systems.

20 Conversely, well targeted ads are relevant and therefore should perform well and enhance a user's experience. Unfortunately, it may be difficult for advertisers to learn how to pick good keyword targets.

25 Selecting good keyword targets may come with a difficult learning curve, and certain problems often occur repeatedly. For example, an advertiser may pick its own keywords, its ads may run for a brief time, and several keywords (or their whole campaign) may get shut off because of low performance (e.g., a low click-through rate). Such advertisers may then request keywords, and recommendations (e.g., from customer service or from an automated keyword recommendation facility) may be provided to the advertiser. However, if the keyword recommendations aren't good, the ad can continue to perform poorly,

30 and get shut off again. This learning process may become a source of advertiser frustration. Such frustration can lead to advertisers to stop using an ad serving system altogether, to reduce their use of an ad serving system, or not to expand

their use of an ad serving system. Even if the advertiser sticks with the ad serving system, a process which slowly converges on a good keyword targeting slows the sales cycle.

Accordingly, there is a need to improve the performance of online advertising. More specifically, there is a need to increase the relevancy of ads, and to exploit opportunities to serve an ad that might otherwise be missed. Doing so could increase the revenue of an ad serving entity, improve the experience of users, and improve efficacy and cost of advertiser's ads.

10 **§ 2. SUMMARY OF THE INVENTION**

The present invention may be used to help advertisers better target the serving of their ads by having more keyword suggestions that are category-aware (and field-proven). For advertisers with new campaigns, the advertiser can enter ad information, such as a creative for example. The keyword facility may use this entered information as seed information to infer one or more categories. It may then request that the advertiser confirm or deny some basic feedback information (e.g., categories, Webpage information, etc.). For example, an advertiser may be provided with candidate categories and may be asked to confirm (e.g., using checkboxes) which of the categories are relevant to their ad. In one embodiment, the ad server system may then automatically provide keywords to the advertiser for its ad.

In one embodiment of the present invention, the ad server system can run a trial on determined keywords. It may then provide recommendations to the advertiser using results of the trial. Some recommendations may be as simple as the recommendation of negative keywords (which will increase the customer's click-through rate, which may save it from being shut off, and/or allow the advertiser to get clicks at a lower CPC while maintaining their rank). Advertisers that have been running ads can request keyword suggestions. The ad serving system could contact such advertisers (e.g., via e-mail) with notices like, "you could get X more clicks at your current CPC if you used {suggested keywords}," or "using {suggested negative keywords} as negative keywords, you could get

your clicks cheaper", or "use {suggested keywords} to avoid shut off," etc.

Keyword suggestions may also be used for building new campaigns and projecting inventory available. For existing campaigns that are not meeting targets, keyword suggestions make it easy to find and add more keywords.

5        Thus, one method consistent with the present invention may include harvesting keyword targets, building category pools, augmenting ad targeting information with information from similar advertisers (e.g., linked by advertiser "category" and keyword intersection), providing category recommendations and accepting advertiser feedback (e.g., decline, accept, confirm, etc.) with a user

10      interface. It may also include running test ads for advertisers on unused inventory, by using category pools or competitor pre-qualified keywords, and suggesting keywords that performed well on the test ads to the end user. In addition to helping advertisers to select targeting keywords, one method consistent with the present invention may also suggest negative category

15      targeting (and negative keyword targeting using the negative categories).

### **§ 3. BRIEF DESCRIPTION OF THE DRAWINGS**

Figure 1 is a high-level diagram showing parties or entities that can interact with an advertising system.

20        Figure 2 is a bubble chart of an exemplary advertising environment in which, or with which, the present invention may operate.

Figure 3 is a bubble chart of operations, consistent with the present invention, that may be used to associate keywords and categories in an online advertising environment, such as the one in Figure 2.

25        Figure 4 is a bubble chart of operations, consistent with the present invention, that may be used to suggest and/or automatically supply keywords for targeting an ad.

Figure 5 is a bubble chart of operations, consistent with the present invention, that may be used to compile unused ad inventory information.

30        Figure 6 is a flow diagram of an exemplary method that may be used to suggest or provide keywords in a manner consistent with the present invention.

Figure 7 is a flow diagram of an exemplary method that may be used to support a user interface in a manner consistent with the present invention.

Figure 8 is a flow diagram of an exemplary method that may be used to test keywords in a manner consistent with the present invention.

5       Figure 9 is a block diagram of apparatus that may be used to effect at least some of the various operations that may be performed consistent with the present invention.

#### **§ 4. DETAILED DESCRIPTION**

10       The present invention may involve novel methods, apparatus, message formats and/or data structures for suggesting and/or providing keywords used to target ads (referred to as “targeting keywords”). The following description is presented to enable one skilled in the art to make and use the invention, and is provided in the context of particular applications and their requirements. Various 15       modifications to the disclosed embodiments will be apparent to those skilled in the art, and the general principles set forth below may be applied to other embodiments and applications. Thus, the present invention is not intended to be limited to the embodiments shown and the inventors regard their invention as any patentable subject matter described.

20       In the following, environments in which, or with which, the present invention may operate are described in § 4.1. Then, exemplary embodiments of the present invention are described in § 4.2. Examples illustrating operations of exemplary embodiments of the present invention are described in § 4.3. Finally, some conclusions regarding the present invention are set forth in § 4.4.

25

##### **§ 4.1 ENVIRONMENTS IN WHICH, OR WITH WHICH, THE PRESENT INVENTION MAY OPERATE**

###### **§ 4.1.1 EXEMPLARY ADVERTISING ENVIRONMENT**

30       Figure 1 is a high level diagram of an advertising environment. The environment may include an ad entry, maintenance and delivery system 120. Advertisers 110 may directly, or indirectly, enter, maintain, and track ad

information in the system 120. The ads may be in the form of graphical ads such as so-called banner ads, text only ads, image ads, audio ads, video ads, ads combining one or more of any of such components, etc. The ads may also include embedded information, such as a link, and/or machine executable 5 instructions. Ad consumers 130 may submit requests for ads to, accept ads responsive to their request from, and provide usage information to, the system 120. Although not shown, other entities may provide usage information (e.g., whether or not a conversion or click-through related to the ad occurred) to the system 120. This usage information may include measured or observed user 10 behavior related to ads that have been served.

One example of an ad consumer 130 is a general content server that receives requests for content (e.g., articles, discussion threads, music, video, graphics, search results, web page listings, etc.), and retrieves the requested content in response to, or otherwise services, the request. The content server 15 may submit a request for ads to the system 120. Such an ad request may include a number of ads desired. The ad request may also include content request information. This information may include the content itself (e.g., page), a category corresponding to the content or the content request (e.g., arts, business, computers, arts-movies, arts-music, etc.), part or all of the content 20 request, content age, content type (e.g., text, graphics, video, audio, mixed media, etc.), geolocation information, etc.

The content server may combine the requested content with one or more of the advertisements provided by the system 120. This combined information including the content and advertisement(s) is then forwarded towards the end 25 user that requested the content, for presentation to the user. Finally, the content server may transmit information about the ads and how, when, and/or where the ads are to be rendered (e.g., position, click-through or not, impression time, impression date, size, conversion or not, etc.) back to the system 120. Alternatively, or in addition, such information may be provided back to the system 30 120 by some other means.

Another example of an ad consumer 130 is a search engine. A search engine may receive queries for search results. In response, the search engine may retrieve relevant search results (e.g., from an index of Web pages). An exemplary search engine is described in the article S. Brin and L. Page, "The 5 Anatomy of a Large-Scale Hypertextual Search Engine," Seventh International World Wide Web Conference, Brisbane, Australia and in U.S. Patent No. 6,285,999 (both incorporated herein by reference). Such search results may include, for example, lists of Web page titles, snippets of text extracted from those Web pages, and hypertext links to those Web pages, and may be grouped 10 into a predetermined number of (e.g., ten) search results.

The search engine may submit a request for ads to the system 120. The request may include a number of ads desired. This number may depend on the search results, the amount of screen or page space occupied by the search results, the size and shape of the ads, etc. In one embodiment, the number of 15 desired ads will be from one to ten, and preferably from three to five. The request for ads may also include the query (as entered or parsed), information based on the query (such as geolocation information, whether the query came from an affiliate and an identifier of such an affiliate), and/or information associated with, or based on, the search results. Such information may include, 20 for example, identifiers related to the search results (e.g., document identifiers or "docIDs"), scores related to the search results (e.g., information retrieval ("IR") scores such as dot products of feature vectors corresponding to a query and a document, Page Rank scores, and/or combinations of IR scores and Page Rank scores), snippets of text extracted from identified documents (e.g., WebPages), 25 full text of identified documents, feature vectors of identified documents, etc.

The search engine may combine the search results with one or more of the advertisements provided by the system 120. This combined information including the search results and advertisement(s) is then forwarded towards the user that requested the content, for presentation to the user. Preferably, the 30 search results are maintained as distinct from the ads, so as not to confuse the user between paid advertisements and presumably neutral search results.

Finally, the search engine may transmit information about the ad and when, where, and/or how the ad was to be rendered (e.g., position, click-through or not, impression time, impression date, size, conversion or not, etc.) back to the system 120. Alternatively, or in addition, such information may be provided back 5 to the system 120 by some other means.

#### **§ 4.1.2 EXEMPLARY AD ENTRY, MAINTENANCE AND DELIVERY ENVIRONMENT**

Figure 2 illustrates an exemplary ad system 120' in which, or with which, 10 the present invention may be used. The exemplary ad system 120' may include an inventory system 210 and may store ad information 205 and usage information 245. The exemplary system 120' may support ad information entry and management operations 215, campaign (e.g., targeting) assistance operations 220, accounting and billing operations 225, ad serving operations 15 230, relevancy determination operations 235, optimization operations 240, relative presentation attribute assignment (e.g., position ordering) operations 250, fraud detection operations 255, and result interface operations 260.

Advertisers 110 may interface with the system 120' via the ad information entry and management operations 215 as indicated by interface 216. Ad 20 consumers 130 may interface with the system 120' via the ad serving operations 230 as indicated by interface 231. Ad consumers 130 and/or other entities (not shown) may also interface with the system 120' via results interface operations 260 as indicated by interface 261.

An advertising program may include information concerning accounts, 25 campaigns, creatives, targeting, etc. The term "account" relates to information for a given advertiser (e.g., a unique email address, a password, billing information, etc.). A "campaign" or "ad campaign" refers to one or more groups of one or more advertisements, and may include a start date, an end date, budget information, geo-targeting information, syndication information, etc. For 30 example, Honda may have one advertising campaign for its automotive line, and a separate advertising campaign for its motorcycle line. The campaign for its

automotive line have one or more ad groups, each containing one or more ads. Each ad group may include a set of keywords, and a maximum cost (cost per click-through, cost per conversion, etc.). Alternatively, or in addition, each ad group may include an average cost (e.g., average cost per click-through, average cost per conversion, etc.). Therefore, a single maximum cost and/or a single average cost may be associated with one or more keywords. As stated, each ad group may have one or more ads or "creatives" (That is, ad content that is ultimately rendered to an end user.). Naturally, the ad information 205 may include more or less information, and may be organized in a number of different ways.

The ad information 205 can be entered and managed via the ad information entry and management operations 215. Campaign (e.g., targeting) assistance operations 220 can be employed to help advertisers 110 generate effective ad campaigns. For example, the campaign assistance operations 220 can use information provided by the inventory system 210, which, in the context of advertising for use with a search engine, may track all possible ad impressions, ad impressions already reserved, and ad impressions available for given keywords. The ad serving operations 230 may service requests for ads from ad consumers 130. The ad serving operations 230 may use relevancy determination operations 235 to determine candidate ads for a given request. The ad serving operations 230 may then use optimization operations 240 to select a final set of one or more of the candidate ads. The ad serving operations 230 may then use relative presentation attribute assignment operations 250 to order the presentation of the ads to be returned. The accounting/billing operations 225 may be used to track charges related to the serving of advertisements and to bill advertisers. The fraud detection operations 255 can be used to reduce fraudulent use of the advertising system (e.g., by advertisers), such as through the use of stolen credit cards. Finally, the results interface operations 260 may be used to accept result information (from the ad consumers 130 or some other entity) about an ad actually served, such as whether or not click-through occurred, whether or not conversion occurred (e.g., whether the

sale of an advertised item or service was initiated or consummated within a predetermined time from the rendering of the ad), etc. Such results information may be accepted at interface 261 and may include information to identify the ad and time the ad was served, as well as the associated result.

5 As will be appreciated from the description to follow, various aspects of the present invention may be applied to campaign targeting assistance operations 220.

#### **S 4.1.3 DEFINITIONS**

10 Online ads, such as those used in the exemplary systems described above with reference to Figures 1 and 2, or any other system, may have various intrinsic features. Such features may be specified by an application and/or an advertiser. These features are referred to as "ad features" below. For example, in the case of a text ad, ad features may include a title line, ad text, and an 15 embedded link. In the case of an image ad, ad features may include images, executable code, and an embedded link. Depending on the type of online ad, ad features may include one or more of the following: text, a link, an audio file, a video file, an image file, executable code, embedded information, etc.

20 When an online ad is served, one or more parameters may be used to describe how, when, and/or where the ad was served. These parameters are referred to as "serving parameters" below. Serving parameters may include, for example, one or more of the following: features of (including information on) a page on which the ad was served, a search query or search results associated with the serving of the ad, a user characteristic (e.g., their geolocation, the 25 language used by the user, the type of browser used, previous page views, previous behavior), a host or affiliate site (e.g., America Online, Google, Yahoo) that initiated the request, an absolute position of the ad on the page on which it was served, a position (spatial or temporal) of the ad relative to other ads served, an absolute size of the ad, a size of the ad relative to other ads, a color of the ad, 30 a number of other ads served, types of other ads served, time of day served,

time of week served, time of year served, etc. Naturally, there are other serving parameters that may be used in the context of the invention.

Although serving parameters may be extrinsic to ad features, they may be associated with an ad as serving conditions or constraints. When used as

5 serving conditions or constraints, such serving parameters are referred to simply as “serving constraints” (or “targeting criteria”). For example, in some systems, an advertiser may be able to target the serving of its ad by specifying that it is only to be served on weekdays, no lower than a certain position, only to users in a certain geolocation, etc. As another example, in some systems, an advertiser  
10 may specify that its ad is to be served only if a page or search query includes certain keywords or phrases. As yet another example, in some systems, an advertiser may specify that its ad is to be served only if a document being served includes certain topics or concepts, or falls under a particular cluster or clusters, or some other classification or classifications.

15 In targeted advertising, a “targeting keyword” may be an identified keyword or group of keywords (including phrases, or negative keywords), that target the serving of the ad by specifying (or restricting) when the ad is eligible to be shown. In some ad serving systems, an advertiser can specify how its keywords are to be used in ad serving decisions by specifying match types. For  
20 example, if an advertiser specifies a “broad” match type (which may be a default setting), the advertiser may include general keyword or keyword phrases (such as “*tennis shoes*” for example) in its keyword list. In this example, the advertiser’s ads are eligible for serving when users search for “*tennis*” and “*shoes*,” in any order, and possibly along with other terms. The advertiser’s ads  
25 might also be subject to expanded matches, including plurals and relevant variations. Since other advertisers may be competing for the same broad-matched keyword combinations that trigger the serving of ads, using broad targeting keywords may increase cost amounts. Using exact, phrase, or negative matches (discussed below) can help advertisers minimize costs. An  
30 advertiser may specify a “phrase” match type (e.g., by entering their keyword in quotation marks, as in “*tennis shoes*”). In this case, its ad will be eligible for

serving when a user searches on the phrase “*tennis shoes*,” in this order, and possibly with other terms in the query. For example, the advertiser’s ad will be eligible for serving for the query “*red tennis shoes*” but not for “*shoes for tennis*.”

Phrase matching is more targeted than broad matching, but slightly more flexible

5 than exact matching. An advertiser may specify an “exact” match type (e.g., by surrounding its keywords in brackets—such as *[tennis shoes]*). In this case, its ads will be eligible for serving when users search for the specific phrase “*tennis shoes*,” in this order, and without any other terms in the query. For example, the advertiser’s ad won’t be eligible for serving for the query “*red tennis shoes*.” An  
10 advertiser may also specify “negative” keywords. For example, if an advertiser’s keyword is *tennis shoes* and it adds the negative keyword *-red*, its ad will not be eligible for serving if a user searches on “*red tennis shoes*.” Naturally, other types of targeting keywords are possible.

“Geolocation information” may include information specifying one or more  
15 of one or more countries, one or more (inter-country) regions, one or more states, one or more metro areas, one or more cities, one or more towns, one or more boroughs, one or more areas with common zip codes, one or more areas with common telephone area codes, one or more areas served by common cable head end stations, one or more areas served by common network access points  
20 or nodes, etc. It may include latitude and/or longitude, or a range thereof. It may include information, such as an IP address, from which a user location can be estimated.

“Ad information” may include any combination of ad features, ad serving constraints, information derivable from ad features or ad serving constraints  
25 (referred to as “ad derived information”), and/or information related to the ad (referred to as “ad related information”), as well as an extension of such information (e.g., information derived from ad related information).

A “document” is to be broadly interpreted to include any machine-readable and machine-storable work product. A document may be a file, a combination of  
30 files, one or more files with embedded links to other files, etc. The files may be of any type, such as text, audio, image, video, etc. Parts of a document to be

rendered to an end user can be thought of as “content” of the document. A document may include “structured data” containing both content (words, pictures, etc.) and some indication of the meaning of that content (for example, e-mail fields and associated data, HTML tags and associated data, etc.) Ad spots in

5 the document may be defined by embedded information or instructions. In the context of the Internet, a common document is a Web page. Web pages often include content and may include embedded information (such as meta information, hyperlinks, etc.) and/or embedded instructions (such as Javascript, etc.). In many cases, a document has a unique, addressable, storage location

10 and can therefore be uniquely identified by this addressable location. A universal resource locator (URL) is a unique address used to access information on the Internet.

“Document information” may include any information included in the document, information derivable from information included in the document

15 (referred to as “document derived information”), and/or information related to the document (referred to as “document related information”), as well as an extensions of such information (e.g., information derived from related information). An example of document derived information is a classification based on textual content of a document. Examples of document related

20 information include document information from other documents with links to the instant document, as well as document information from other documents to which the instant document links.

Content from a document may be rendered on a “content rendering application or device”. Examples of content rendering applications include an

25 Internet browser (e.g., Explorer or Netscape), a media player (e.g., an MP3 player, a Realnetworks streaming audio file player, etc.), a viewer (e.g., an Adobe Acrobat pdf reader), etc.

A “content owner” is a person or entity that has some property right in the content of a document. A content owner may be an author of the content. In

30 addition, or alternatively, a content owner may have rights to reproduce the content, rights to prepare derivative works of the content, rights to display or

perform the content publicly, and/or other proscribed rights in the content.

Although a content server might be a content owner in the content of the documents it serves, this is not necessary.

“User information” may include user behavior information and/or user 5 profile information. It may also include a user’s geolocation, or an estimation of the user’s geolocation.

“E-mail information” may include any information included in an e-mail (also referred to as “internal e-mail information”), information derivable from 10 information included in the e-mail and/or information related to the e-mail, as well as extensions of such information (e.g., information derived from related information). An example of information derived from e-mail information is information extracted or otherwise derived from search results returned in response to a search query composed of terms extracted from an e-mail subject line. Examples of information related to e-mail information include e-mail

15 information about one or more other e-mails sent by the same sender of a given e-mail, or user information about an e-mail recipient. Information derived from or related to e-mail information may be referred to as “external e-mail information.”

Various exemplary embodiments of the present invention are now 20 described in § 4.2.

## **§ 4.2 EXEMPLARY EMBODIMENTS**

As indicated above, the present invention may be used to suggest and/or provide keywords used to target ads. Exemplary embodiments for associating 25 keywords and categories are described in § 4.2.1 below. Then, exemplary embodiments for using such associations to suggest and/or provide keywords is described in § 4.2.2. Some possible refinements and alternative embodiments are discussed in § 4.2.3. Finally, exemplary apparatus that may be used to effect various aspects of the invention are described in § 4.2.4.

### **30 § 4.2.1 ASSOCIATING KEYWORDS AND CATEGORIES**

Figure 3 is a bubble chart of operations that may be performed, and information that may be used or stored, consistent with the present invention.

More specifically, the operations of Figure 3 may be used to generate associations between keywords and categories. For example, keyword

- 5 classification operations 320 may be used to generate a keyword index 330 from keywords 310. The keywords 310 may be from search queries for example. The keyword index 330 may use a keyword 332 as a lookup key to one or more associated categories 334. Sorting operations 360 may be used to sort the entries of the keyword index 330. For example, the entries may be sorted based
- 10 on frequency of use of keywords 332. Finally, index inversion operations 340 may be used to generate a keyword inverted index 350 from the keyword index 330. The keyword inverted index 350 may use a category 352 as a lookup key to associated keywords 354.

15 **§ 4.2.2 USING KEYWORD AND CATEGORY ASSOCIATIONS TO SUGGEST AND/OR PROVIDE KEYWORDS**

Figure 4 is a bubble chart of operations, consistent with the present invention, that may be used to suggest and/or automatically supply keywords for targeting an ad. Category determination operations 410 may use ad information 420 to determine one or more categories 450. The category determination operations 410 may determine one or more categories using one or more of the following techniques. For example, using one technique, if the ad information 420 includes one or more keywords, the category determination operations may 25 use keyword-category association information, such as a keyword index 330, to determine one or more categories 334 associated with the one or more keywords. The one or more categories determined may be provided to user interface operations 440. The user interface operations 440 may be used to (i) present determined categories to an advertiser, or its representative (simply referred to as “advertiser”), and to receive feedback concerning the presented categories from the advertiser. For example, a number of candidate categories 30 may be presented to an advertiser. In response, the advertiser can indicate

which of the candidate categories are appropriate and/or inappropriate. In another technique, the ad information 410 may include a creative or a link to a landing Webpage. The category determination operations 410 may use such information, as well as other information 405, to determine one or more

5 categories. (See, for example, U.S. Patent Application Serial No. 10/419,692 (incorporated herein by reference), titled "DETERMINING CONTEXTUAL INFORMATION FOR ADVERTISEMENTS AND USING SUCH DETERMINED CONTEXTUAL INFORMATION TO SUGGEST TARGETING CRITERIA AND/OR IN THE SERVING OF ADVERTISEMENTS," filed on April 21, 2003 and  
10 listing Amit Singhal, Mehran Sahami, Amit Patel and Steve Lawrence as inventors. That application stated that generated ad context information may be used to help advertisers consider additional serving constraints, such as keywords for example.)

Regardless of the technique or techniques used, one or more categories  
15 450 are generated. Keyword suggestion/provision operations 460 may use the one or more generated categories 450, to generate one or more keywords 465. For example, the keyword suggestion/provision operations 460 may use the one or more categories 450 as keys 352 to lookup keywords 354 in a keyword inverted index 350. The determined keywords 476 may be provided to the user  
20 as suggested keywords, or may automatically be associated with an ad as targeting keywords. However, in one embodiment of the present invention described below, the keywords are tested for qualification first.

In one embodiment of the present invention, the keyword suggestion/provision operations 460 may use unused inventory information 455 in its determination of keywords. For example, such unused inventory information 455 may associate keywords with a number of unused ad spots over a period of time. Generated targeting keywords 465 may be ordered based on number of unused ad spots associated with the keywords. In this way, keywords that, if used as targeting keywords, would fill many otherwise unused ad spots  
30 may be preferred over those that would fill few otherwise unused ad spots.

Referring to Figure 5, unused inventory information 560 may be compiled by page assembly operations 510. These operations 510 basically mix Web page information (e.g., search results) 530, query information 520 and ads (e.g., from an ad server) 540 to generate a Web page with ads 550. Suppose that the

5 Web page has ten (10) ad spots and ten (10) ads are served. In this case, there are no unused ad spots, and the information 560 need not be updated. If, however, the Web page has ten (10) ad spots and only three (3) ads are served, there are seven (7) unused ad spots. In this case, entries of the information 560 corresponding to keywords 562 found in the query 520 may have an associated  
10 number of unused ad spots 564 increased. Thus, the number of unused ad spots 564 may correspond to a sum of unused ad spots over a number of page serves. The number 564 may be a running total, may have a time windowing function applied (e.g., based on page serves for the past two weeks), etc. The unused inventory information 560 may be ordered from time to time based on the  
15 number of unused ad spots.

Referring back to Figure 4, as described above, although the determined keywords can simply be suggested to the advertiser, or provided to the ad as keyword targets, it may be desirable to test the keywords for qualification first. In an embodiment in which keywords are tested for qualification, keyword status

20 check operations 470 may check keyword status information (e.g., one or more of: qualified for use as ad targeting information without need for advertiser approval; qualified for use as ad targeting information but only with advertiser approval; neither qualified, nor disqualified; or disqualified for use as ad targeting) information of the keyword(s) 465 under consideration. The keyword  
25 status information 475 may be maintained for a keyword in general (e.g., without regard to a category of the ad), and/or may be maintained for a keyword in the context of one or more particular categories. Although not shown in Figure 4, the keyword status information 475 may be stored in, or otherwise linked with, the keyword inverted index 350. If the keyword is qualified for use as ad targeting  
30 information without the need for advertiser approval, it may be added to the ad information as a targeting keyword. If the keyword is qualified for use as ad

targeting information but only with advertiser approval, it may be provided to user interface operations 440 for presentation to the advertiser. If the keyword is disqualified, it should be dropped from consideration. Finally, if the keyword is neither qualified, nor disqualified (e.g., if there is no qualification information for

5 the keyword), it may be passed to keyword trial operations 480.

The keyword trial operations 480 may use the keyword as a targeting keyword to serve an ad 485 to see how it will perform. The trial may be limited to otherwise unused inventory. Generally, if the ad performs well, it may be qualified, if it performs poorly, it may be disqualified, and the keyword status

10 information 475 may be updated appropriately. Various status levels may be associated with various levels of trial performance. Other ways of qualifying or disqualifying a keyword may be used. If there are a number of keywords that are under consideration but that don't have qualification information, keyword trial operations 480 may try the keywords with the highest amount of unused

15 inventory first. Then additional keywords may be tried. In one embodiment of the present invention, the additional keywords may be from those with more unused inventory to those with less unused inventory.

Having introduced operations that may be performed in accordance with the present invention, a number of exemplary methods for performing the

20 operations are now described.

Figure 6 is a flow diagram of an exemplary method 600 that may be used to perform keyword suggestion/provision operations 460 in a manner consistent with the present invention. One or more categories are accepted. (Block 610) One or more keywords are determined using the accepted categories. (Block

25 620) Recall that this may be done using a keyword index 330. The determined keywords may be ordered using unused inventory information. (Block 630) In this way, keywords that, if used as targeting keywords, would fill more ad spots may be considered first. As indicated by loop 640-660, a number of acts may be performed for each of the determined keywords, or at least of the determined

30 keywords under consideration. In one embodiment of the present invention, the keyword is simply suggested to the advertiser or provided as a targeting

keyword. (Block 650) In an alternative embodiment of the present invention, keyword status is checked. For example, it may be determined whether or not the keyword is qualified. (Decision block 680) If so, the method 600 continues to block 650 already discussed. If the keyword is not qualified (and hasn't been 5 disqualified), the password may be passed to trial operations for qualification testing. (Block 690) Once all of the keywords have been processed, the method 600 is left. (Node 670)

Figure 7 is a flow diagram of an exemplary method 700 that may be used to perform user interface operations 440 in a manner consistent with the present

10 invention. As indicated by block 710, various branches of the method 700 may be performed in response to the occurrence of various different events. More specifically, if one or more candidate categories are provided (e.g., by category determination operations 410), the candidate category(ies) may be presented to the advertiser to accept or decline (Block 720) before the method 700 branches 15 back to block 710. If one or more candidate keyword(s) are provided (e.g., by keyword suggestion operations 460), the candidate keyword(s) may be presented to the advertiser to accept or decline (Block 730) before the method 700 branches back to block 710. If one or more categories are accepted by the advertiser, the category(ies) may be passed to keyword suggestion/provision 20 operations, such as those 460 of Figure 4 (Block 740) before the method 700 branches back to block 710. If one or more keyword(s) are accepted by the advertiser, the keyword (s) may be added to ad information as a targeting keyword (Block 750) before the method 700 branches back to block 710. If an exit condition occurs, the method 700 may leave. (Node 760)

25 Figure 8 is a flow diagram of an exemplary method 800 that may be used to try keywords for qualification as targeting keywords in a manner consistent with the present invention. A keyword (or more than one keyword) is accepted.

(Block 810). The trial operations may use one or more keywords as targeting keywords in the serving of an ad (or even a group of ads) (Block 820) and the 30 performance of such ads may be tracked (Block 830). In one embodiment of the present invention, the serving of the ads using trial targeting keyword (s) may be

limited to ad spots (inventory) that otherwise would be unused. After a certain amount of time and/or after a certain number of such ad serves, various branches of the method 800 may be performed responsive to various different performance levels. If a keyword performs well (e.g., in general, or for a 5 particular category), it may be marked as a keyword to be suggested, and/or as a qualified keyword (e.g., in general, or for the particular category) (Block 850) before the method 800 is left (Node 870). In fact, generic creatives (for example, creative templates with certain advertiser information inserted) with qualified targeting keywords could be subject to expedited approvals. If, on the other 10 hand, a keyword does not perform well (e.g., in general, or for a particular category), it may be disqualified and marked as unusable (e.g., in general, or for the particular category) (Block 860) before the method 800 is left (Node 870). In this way, advertisers can avoid the frustration of targeting their ads using poorly performing keywords. Although not shown, keywords can be assigned various 15 different status levels associated with various different levels of performance. In one embodiment, performance for one or more keywords may be considered to be good if ads served pursuant to using the keyword(s) as targeting keywords perform (e.g., have a click-through rate) comparable to what salespeople and/or customers already think are the best keywords (e.g., the keywords that they are 20 already using).

#### **§ 4.2.3 REFINEMENTS AND ALTERNATIVES**

It may be desirable for the ad serving entity to provide advertisers with opaque ad performance reporting (e.g. not by specific keyword targets). In this 25 way, it may become possible to protect a keyword targeting knowledge base and treat the data as a trade secret. In such an embodiment, advertisers may be satisfied to know that their ad campaign is performing well, without needing to know the details about keyword targeting used. This may make alternative keyword targeting ad serving firms less attractive to an advertiser.

30 Note that qualified keywords may be any type of keyword, such as negative keywords for example.

#### **§ 4.2.4 EXEMPLARY APPARATUS**

Figure 9 is high-level block diagram of a machine 900 that may perform one or more of the operations and store various information discussed above.

- 5 The machine 900 basically includes a processor(s) 910, an input/output interface unit(s) 930, a storage device(s) 920, and a system bus or network 940 for facilitating the communication of information among the coupled elements. An input device(s) 932 and an output device(s) 934 may be coupled with the input/output interface(s) 930.
- 10 The processor(s) 910 may execute machine-executable instructions (e.g., C or C++ running on the Solaris operating system available from Sun Microsystems Inc. of Palo Alto, California or the Linux operating system widely available from a number of vendors such as Red Hat, Inc. of Durham, North Carolina) to perform one or more aspects of the present invention. At least a portion of the machine executable instructions may be stored (temporarily or more permanently) on the storage device(s) 920 and/or may be received from an external source via an input interface unit 930.
- 15

In one embodiment, the machine 900 may be one or more conventional personal computers. In this case, the processing unit(s) 910 may be one or more microprocessors. The bus 940 may include a system bus. The storage devices 920 may include system memory, such as read only memory (ROM) and/or random access memory (RAM). The storage device(s) 920 may also include a hard disk drive for reading from and writing to a hard disk, a magnetic disk drive for reading from or writing to a (e.g., removable) magnetic disk, and an optical disk drive for reading from or writing to a removable (magneto-) optical disk such as a compact disk or other (magneto-) optical media.

A user may enter commands and information into the personal computer through input devices 932, such as a keyboard and pointing device (e.g., a mouse) for example. Other input devices such as a microphone, a joystick, a game pad, a satellite dish, a scanner, or the like, may also (or alternatively) be included. These and other input devices are often connected to the processing

unit(s) 910 through an appropriate interface 930 coupled to the system bus 940.

The output device(s) 934 may include a monitor or other type of display device, which may also be connected to the system bus 940 via an appropriate interface. In addition to (or instead of) the monitor, the personal computer may

5 include other (peripheral) output devices (not shown), such as speakers and printers for example.

#### **§ 4.3 EXAMPLES OF OPERATIONS**

An example illustrating exemplary operations of the present invention is

10 now provided. Suppose an advertiser is a Jaguar car dealer. The word “jaguar” may be provided as ad information 420. It may have been taken from entered creative information, information from a landing Webpage specified in the ad, and/or an entered targeting keyword, etc. Category determinations operations 410 may determine various, possibly relevant, categories (and possibly 15 sub-categories) such as:

*automobiles ...*

*computers ... operating systems ...*

*music ... popular music ...*

*music ... musical instruments ...*

*animals ... mammals ... felines ...*

*movies ... foreign films ...*

*travel ... resorts ...*

*sports & recreation ... snorkeling ... scuba ...*

*sports & recreation ... football ...*

*pets ... fish*

For example, “jaguar” may be relevant to the luxury automobile, the apple Macintosh operating system, the recording artist (Jaguar Wright or Jet Jaguar),

30 the foreign movie (Aimee and Jaguar), the reef in Belize, the football team, or the cichlid fish. The determined categories may be provided to the advertiser via user interface operations 440 and the advertiser may specify that only the automobile category is relevant. The “automobile” category may then be

provided to keyword suggestion/provision operations 460 which may generate one or more keywords, such as:

***cars, parts, service, dealer, loaner, X-type, S-type, XJ ...***

5

therefrom. The generated keywords may (a) be added to the ad as targeting keywords, (b) suggested to the advertiser as targeting keywords, or (c) tested if not already qualified. Keyword trial operations 480 may run the ad with the keyword(s) under consideration and track the performance of the ad with respect

10 to one or more of the keyword(s) under consideration. Keywords that perform well may (a) be added to the ad as targeting keywords, or (b) suggested to the advertiser as targeting keywords. Keyword that do not perform well may be disqualified in general, or disqualified for the category automobile.

15           **§ 4.4 CONCLUSIONS**

As can be appreciated from the foregoing disclosure, the present invention can be used to help online advertisers and entities serving online ads by (i) suggesting or providing targeting criteria to better target the ad and to exploit otherwise missed opportunities to serve relevant ads, and (ii) helping advertisers

20 to eliminate poorly targeted ads. The various techniques described above may be used in combination or in concert. Some embodiments of the present invention unburden advertisers of the need to learn the details of keyword targeting. Moreover, valuable information is learned from queries where few ads are otherwise served. The present invention can improve the targeting of ads, 25 improve customer experience, reduced load on customer support, and increase revenue.